

Submitted

Tuning Parameter Calibration in High-dimensional Logistic Regression With Theoretical Guarantees

BY WEI LI

School of Mathematical Sciences, Peking University, Beijing, China

weylpeking@pku.edu.cn

JOHANNES LEDERER

Departments of Statistics and Biostatistics, University of Washington, Seattle, WA, USA

ledererj@uw.edu

SUMMARY

Feature selection is a standard approach to understanding and modeling high-dimensional classification data, but the corresponding statistical methods hinge on tuning parameters that are difficult to calibrate. In particular, existing calibration schemes in the logistic regression framework lack any finite sample guarantees. In this paper, we introduce a novel calibration scheme for penalized logistic regression. It is based on simple tests along the tuning parameter path and satisfies optimal finite sample bounds. It is also amenable to easy and efficient implementations, and it rivals or outmatches existing methods in simulations and real data applications.

Some key words: Feature selection; Penalized logistic regression; Tuning parameter calibration

1. INTRODUCTION

The advent of high-throughput technology has created a large demand for feature selection with high-dimensional classification data. In gene expression analysis or genome-wide association studies, for example, investigators attempt to select from a large set of potential risk factors the predictors that are most useful in discriminating two or more conditions of interest. The standard approaches for such tasks are penalized likelihood methods (Bühlmann & van de Geer, 2011; Hastie et al., 2015; Bunea, 2008; Ravikumar et al., 2010; Ryali et al., 2010; Wu et al., 2009). However, the performance of these methods hinges on the calibration of tuning parameters that balance model fit and model complexity.

The focus of this paper is the calibration of the ℓ_1 -penalized likelihood for feature selection in logistic regression. The most widely used schemes for this calibration are based on Cross-Validation (CV) or on information criteria, including the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). However, CV-based procedures can be inconsistent for model selection (Shao, 1993) and computationally expensive, and in any case, they are not

equipped with finite sample bounds. Also AIC- or BIC-type procedures, as well as more recent methods (Sabourin et al., 2015; Zhang et al., 2010; Fan & Tang, 2013), lack in non-asymptotic theory.

In this paper, we introduce a novel calibration scheme. The approach is based on the AV-testing idea introduced by Chichignoud et al. (2014) and on sharp ℓ_∞ -bounds. It is easy to implement and computationally efficient, and in contrast to previous approaches, it is equipped with finite sample guarantees.

The remainder of this paper is organized as follows. Section 2 contains our main proposal and the theoretical results. Section 3 and Section 4 demonstrate that our method is also a contender in simulations and real data applications. Section 5 contains a brief discussion. The proofs and further simulations are deferred to the Appendix.

Notation. The index sets are denoted by $[k] = \{1, \dots, k\}$ for $k \in \{1, 2, \dots\}$, and the cardinality of sets is denoted by $|\cdot|$. For a given vector $\beta \in \mathbb{R}^p$, the support set of β is written as $\text{supp}(\beta) = \{j \in [p] : \beta_j \neq 0\}$, and for $q \in [1, \infty]$, the ℓ_q -norm of β is denoted by $\|\beta\|_q$. The ℓ_q -induced matrix-operator norms are denoted by $\|\cdot\|_q$. Two examples are the spectral norm $\|\cdot\|_2$, which denotes the maximal singular value of a matrix, and the ℓ_∞ -matrix norm $\|X\|_\infty = \max_{i=1, \dots, n} \sum_{j=1}^p |X_{ij}|$. The minimal and maximal eigenvalue of a square matrix are denoted by $\Omega_{\min}(\cdot)$ and $\Omega_{\max}(\cdot)$, respectively. For a given subset A of $[p]$, the vectors $\beta_A \in \mathbb{R}^{|A|}$ and $\beta_{A^c} \in \mathbb{R}^{|A^c|}$ denote the components of β in A and in its complement A^c , respectively, and given a matrix $X \in \mathbb{R}^{n \times p}$, the matrix X_A denotes the sub-matrix of X with column indexes restricted to A . The diagonal matrix with diagonal elements a_1, \dots, a_n is denoted by $\text{diag}\{a_1, \dots, a_n\}$. The function w is finally defined as $w(u, v) = \exp(u^\top v) / (1 + \exp(u^\top v))^2$ for vectors u, v of the same length.

2. METHODOLOGY

2.1. Model and Assumptions

In this section, we formulate the general setting and introduce the assumptions required for the theoretical analysis. We consider data in the form of a real-valued $n \times p$ design matrix X and a binary response vector $Y = (y_1, \dots, y_n)^\top$. Our framework allows for high-dimensional data, where p rivals or outmatches n . We denote the rows of X (i.e. the samples) by $x_1, \dots, x_n \in \mathbb{R}^p$ and the columns of X (i.e. the predictors) by $x^1, \dots, x^p \in \mathbb{R}^n$. The matrix X can be deterministic or a realization of a random matrix; we only assume that the absolute values of the entries in X are bounded by a common constant $c_b > 0$.

The design matrix X and the response vector Y are linked by the standard logistic regression model

$$\text{pr}(y_i = 1 \mid x_i) = \frac{\exp(x_i^\top \beta^*)}{1 + \exp(x_i^\top \beta^*)} \quad (i = 1, \dots, n), \quad (1)$$

where $\beta^* \in \mathbb{R}^p$ is the unknown regression vector. Our goal is feature selection (or also called support recovery), that is, estimation of the support set $S = \text{supp}(\beta^*)$. The basis for approaching

this task is the well-known family of estimators

$$\hat{\beta}_\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \{L(\beta) + \lambda \|\beta\|_1\} \quad (\lambda > 0) \quad (2)$$

indexed by the tuning parameter λ . The first term $L(\beta) = \sum_{i=1}^n (\log(1 + \exp(x_i^\top \beta)) - y_i x_i^\top \beta)/n$ is the negative log-likelihood function, and the second term is a regularization that exploits that $s = |S| \ll n, p$ in many applications. We estimate S by $\text{supp}(\hat{\beta}_\lambda)$ for a data-driven tuning parameter $\lambda \equiv \lambda(Y, X)$.

Support recovery is feasible only if the correlations in the design matrix X are sufficiently small. In the following, we state corresponding assumptions that virtually coincide with those ones used by Ravikumar et al. (2010) in the context of Ising models. The assumptions are formulated in terms of $W = \text{diag}\{w(x_1, \beta^*), \dots, w(x_n, \beta^*)\}$, the Hessian of the log-likelihood function evaluated at the true regression parameter β^* . We first require that the submatrix of the Hessian matrix corresponding to the relevant covariates has eigenvalues bounded away from zero.

Assumption 1 (Minimal eigenvalue condition). It holds that

$$c_{\min} = \Omega_{\min}(X_S^\top W X_S/n) > 0.$$

Note that if this assumption were violated, the relevant covariates would be linearly dependent, and the true support set S would not be well-defined. Note also that $c_{\min} \rightarrow 1$ for orthogonal design and $\beta^* \rightarrow \mathbf{0}$. Additionally, we set $c_{\max} = \Omega_{\max}(X^\top X/n)$; it holds that $c_{\max} = 1$ for orthogonal design.

We finally impose an irrepresentability condition.

Assumption 2 (Irrepresentability condition). It holds that

$$\gamma = 1 - \|(X_S^\top W X_S)^{-1} X_S^\top W X_{S^c}\|_\infty > 0.$$

This assumption is a modified version of the irrepresentability condition commonly used in the theory for linear regression with the Lasso (Zhao & Yu, 2006). It holds that $\gamma \rightarrow 1$ for orthogonal design and $\beta^* \rightarrow \mathbf{0}$. More generally, irrepresentability conditions prevent the relevant covariates from being strongly correlated with the irrelevant covariates. This ensures that the true support set can be identified with finitely many samples.

For ease of notation, we also set

$$a = \|(X_S^\top W X_S)^{-1}\|_\infty / \|(X_S^\top W X_S)^{-1}\|_2.$$

Again, $a \rightarrow 1$ for orthogonal design and $\beta^* \rightarrow \mathbf{0}$.

Importantly, the above assumptions on the design are not needed in the analysis of the proposed scheme itself. Instead, the assumptions are needed to ensure that there is a viable estimator in the family (2) at all. We discuss this in the following section.

2.2. ℓ_∞ -estimation and Support Recovery

ℓ_∞ -estimation and support recovery are two closely related aspects of high-dimensional logistic regression. In this section, we thus establish oracle inequalities for both these tasks.

To state the result, we define the vector of residuals as $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^\top$ with entries $\varepsilon_i = y_i - \text{pr}(y_i = 1 \mid x_i)$ for $i \in [n]$. The vector ε is random noise with mean zero. We also define the

events

$$\mathcal{T}_\lambda = \left\{ \frac{4(2-\gamma)}{n\gamma} \|X^\top \varepsilon\|_\infty \leq \lambda \right\} \quad (\lambda > 0).$$

For ease of presentation, we assume $\lambda \leq \gamma c_{\min}^2 / (100c_b(2-\gamma)sc_{\max})$ in the remainder.¹ We then find the following result.

THEOREM 1 (ℓ_∞ -BOUND AND SUPPORT RECOVERY). *Under Assumptions 1 and 2, the following properties hold on the event \mathcal{T}_λ .*

- (a) ℓ_∞ -bound: $\|\hat{\beta}_\lambda - \beta^*\|_\infty \leq 1.5a\lambda/c_{\min}$;
- (b) support recovery: $\text{supp}(\hat{\beta}_\lambda) \subset S$, and $\text{supp}(\hat{\beta}_\lambda) = S$ if $\min_{j \in S} |\beta_j^*| > 1.5a\lambda/c_{\min}$.

Oracle inequalities are the standard way to state finite sample bounds in high-dimensional statistics (Bühlmann & van de Geer, 2011). Similar results for ℓ_1 -penalized logistic/linear regression have also been derived elsewhere (Bunea, 2008; Chichignoud et al., 2014; Lounici, 2008; Ravikumar et al., 2010), but the above formulation is particularly useful for our purposes. Part (a) implies that for a suitable tuning parameter λ , the estimator $\hat{\beta}_\lambda$ is uniformly close to the regression vector β^* . Part (b) implies that for suitable tuning parameter, the estimator $\text{supp}(\hat{\beta}_\lambda)$ provides exact support recovery if the non-zero parameters are sufficiently large. As long as the design assumptions are met, Theorem 1 thus ensures that the family (2) contains a viable estimator.

2.3. Testing-based Calibration

Theorem 1 ensures that the family (2) contains an accurate estimator. This leaves us with two tasks: (i) We have to formulate a notation of optimality within the family (2). In other words, we have to define what an optimal tuning parameter is. (ii) We have to formulate a scheme to find an optimal tuning parameter from data. We address these two tasks in the following.

Let us first define the concept of oracle tuning parameters. Since one can handle only finitely many values in practice, we consider a fixed but arbitrary sequence $0 < \lambda_1 < \dots < \lambda_N$ of tuning parameters and denote the corresponding set by $\Lambda = \{\lambda_1, \dots, \lambda_N\}$. In view of Theorem 1, an optimal tuning parameter satisfies two requirements. On the one hand, the bounds hold only on the event \mathcal{T}_λ . Thus, an optimal tuning parameter needs to ensure that the event \mathcal{T}_λ holds with high probability. On the other hand, the bounds are linear in λ . Thus, an optimal tuning parameter should be as small as possible. We formalize this notion as follows:

DEFINITION 1 (ORACLE TUNING PARAMETER). *Given $\delta \in (0, 1)$, the oracle tuning parameter is*

$$\lambda_\delta^* = \underset{\lambda \in \Lambda}{\operatorname{argmin}} \{pr(\mathcal{T}_\lambda) \geq 1 - \delta\}. \quad (3)$$

Since the set Λ is finite, the oracle tuning parameter is always well-defined.

We call the optimal tuning parameter “oracle tuning parameter” to signify that it is a purely theoretical quantity and cannot be used in applications. First, λ_δ^* depends on γ , which is unknown

¹ On a high level, $\gamma c_{\min}^2 / (100c_b(2-\gamma)sc_{\max}) \sim 1$ and $4(2-\gamma)\|X^\top \varepsilon\|_\infty / (n\gamma) \sim 1/\sqrt{n}$. Thus, $\gamma c_{\min}^2 / (100c_b(2-\gamma)sc_{\max}) \gg 4(2-\gamma)\|X^\top \varepsilon\|_\infty / (n\gamma)$. Since the right-hand side of this relation is basically the optimal tuning parameter targeted in our study, see the next section, the much larger upper bound on λ has no impact on our analysis.

in practice. Second, even if γ were known, a precise evaluation of λ_δ^* would be computationally intensive. Finally, it is unclear how to choose δ . We thus aim at finding a data-driven selection rule that mimics the performance of the optimal tuning parameter. The following tests based on the AV-idea (Chichignoud et al., 2014) provide this.

DEFINITION 2 (TESTING-BASED CALIBRATION). *Given a constant $C \geq 1.5a/c_{\min}$, we select the tuning parameter*

$$\hat{\lambda} = \min \left\{ \lambda \in \Lambda \mid \max_{\lambda', \lambda'' \in \Lambda: \lambda', \lambda'' \geq \lambda} \left(\frac{\|\hat{\beta}_{\lambda'} - \hat{\beta}_{\lambda''}\|_\infty}{\lambda' + \lambda''} - C \right) \leq 0 \right\} \quad (4)$$

and set

$$\hat{S} = \{j \in [p] : |(\hat{\beta}_{\hat{\lambda}})_j| \geq 3C\hat{\lambda}\}. \quad (5)$$

Two features are apparent immediately: First, the method is computationally efficient, because it requires at most one pass of the tuning parameter path. This path can be computed by standard algorithms such as `glmnet` (Friedman et al., 2016). Second, the method is easy to implement because it consists of simple ℓ_∞ -tests along the tuning parameter path. The tests also highlight the close connections between ℓ_∞ -estimation and our final goal, support recovery.

The third feature of our scheme is that it is equipped with optimal finite sample theoretical guarantees. We establish this in the following result.

THEOREM 2 (OPTIMALITY OF THE TESTING-BASED CALIBRATION). *Under Assumptions 1 and 2, for any $\delta \in (0, 1)$ and $C \geq 1.5a/c_{\min}$, the tuning parameter $\hat{\lambda}$ from (4) provides with probability at least $1 - \delta$*

$$\hat{\lambda} \leq \lambda_\delta^* \quad \text{and} \quad \|\hat{\beta}_{\hat{\lambda}} - \beta^*\|_\infty \leq 3C\lambda_\delta^*,$$

and, if $\min_{j \in S} |\beta_j^*| > 6C\lambda_\delta^*$,

$$\hat{S} \supset S.$$

Let us highlight some aspects of this result: First, all results are stated for fixed n, p , and all constants are specified. The bounds are thus finite sample bounds that can provide, as opposed to asymptotic bounds, concrete insights into the practical performance of the method. Next, the guarantees hold for any γ and δ , but these quantities do *not* need to be specified in the method. Similarly, the results hold irrespective of the set Λ , in particular, irrespective of the number of tuning parameters N . The set Λ enters the results only through λ_δ^* : the finer the grid Λ , the more precise the optimal tuning parameter λ_δ^* , and thus, the sharper the guarantees. Furthermore, the ℓ_∞ -bounds demonstrate the optimality of the method. Indeed, the estimator with optimal, in practice unknown tuning parameter satisfies $\|\hat{\beta}_{\lambda_\delta^*} - \beta^*\|_\infty \leq 1.5a\lambda_\delta^*/c_{\min}$, see Theorem 1. The bound for the estimator with the data-driven tuning parameter $\hat{\lambda}$ equals this bound - up to a constant factor 3. Finally, since Definition (5) contains a threshold, which is based on the guarantee $\hat{\lambda} \leq \lambda_\delta^*$, the number of false positives is typically small. Yet, the second part of the theorem ensures that \hat{S} contains all sufficiently large predictors, which means that also the number of false negatives is typically small. Theorem 2 thus provides accurate feature selection guarantees for the testing-scheme. We are not aware of any comparable feature selection (or ℓ_∞ -) guarantees for standard calibration schemes.

To summarize, the proposed testing-based method accurately mimics the performance of the optimal tuning parameter, and yet, it is computationally efficient and does not depend on the quantities γ and δ . The only parameter that enters the method is C . The bounds highlight the role of C from a theoretical point of view. However, C can be set to a universal constant for all practical purposes; in particular, it does not need to be calibrated. Note first that the optimal value is $C = 1.5a/c_{\min}$ in view of the theoretical bounds. As described above, support recovery is not possible in highly correlated settings, and it has been pointed out that large β^* can be problematic more generally in penalized estimation (Dalalyan et al., 2016). Therefore, it makes sense to assume near orthogonal design and β^* not too large, so that $a, c_{\min} \approx 1$ and consequently, $1.5a/c_{\min} \approx 1.5$. We thus set $C = 1.5$. The simulation results below indicate that indeed no further calibration is required. The testing-based scheme is thus a practical scheme with a sound theoretical foundation.

3. SIMULATION STUDY

In this section, we show the practical performance of the proposed scheme in a simulation study. We simulate data from the logistic regression model (1) with $n = 200$ samples and $p \in \{200, 500, 2000\}$ predictors. The row vectors x_i of the design matrix X are i.i.d. Gaussian with mean zero and covariance $\Sigma = (1 - \kappa)\mathbf{I} + \kappa\mathbf{1}\mathbf{1}^T$, where \mathbf{I} is the identity matrix, $\mathbf{1}$ the matrix of all ones, and $\kappa \in \{0, 0.25, 0.5, 0.75\}$ the level of correlations among the predictors. The coordinates of the regression vector β^* are set to zero except for $s \in \{5, 10\}$ uniformly at random chosen entries that are set to 1 or -1 with equal probability. We consider $N = 500$ tuning parameters that are equally spaced on $[\lambda_1, \lambda_N]$, where $\lambda_1 = 0.0001\lambda_N$ and $\lambda_N = 10\log(p)/n$. For each of the total 24 settings, we report the means over 200 replications. The methods under consideration are the testing-based method defined in (4) and (5), BIC, 10-fold CV, and AIC. No thresholding is applied for the standard methods, since there is no guidance on the choice of such a threshold. All computations are conducted with the software R (R Core Team, 2016) and the `glmnet` package.

Since our goal is support recovery, we compare the methods in terms of Hamming distance, which is the sum of the number of false positives and false negatives. Figure 1 contains the results for $\kappa = 0.5$ and $\kappa = 0.75$. The results for smaller correlations are deferred to the Appendix. The results allow for two observations: First, BIC consistently outperforms CV and AIC. This is no surprise, given that BIC is specifically designed for feature selection. Second, our testing-based scheme rivals BIC across all settings.

BIC and AIC require one complete pass of the tuning parameter path. 10-fold CV requires one complete pass of 10 tuning parameter paths and thus, requires about 10 times more computational power (or parallelization). The testing-based scheme is the most efficient approach: it requires at most one complete pass of the tuning parameter path, and typically even less, because it stops as soon as the tuning parameter is selected. For illustration, Figure 2 summarizes the run times for four settings with $\kappa \in \{0.5, 0.75\}$; the results with smaller correlations are deferred to the Appendix.

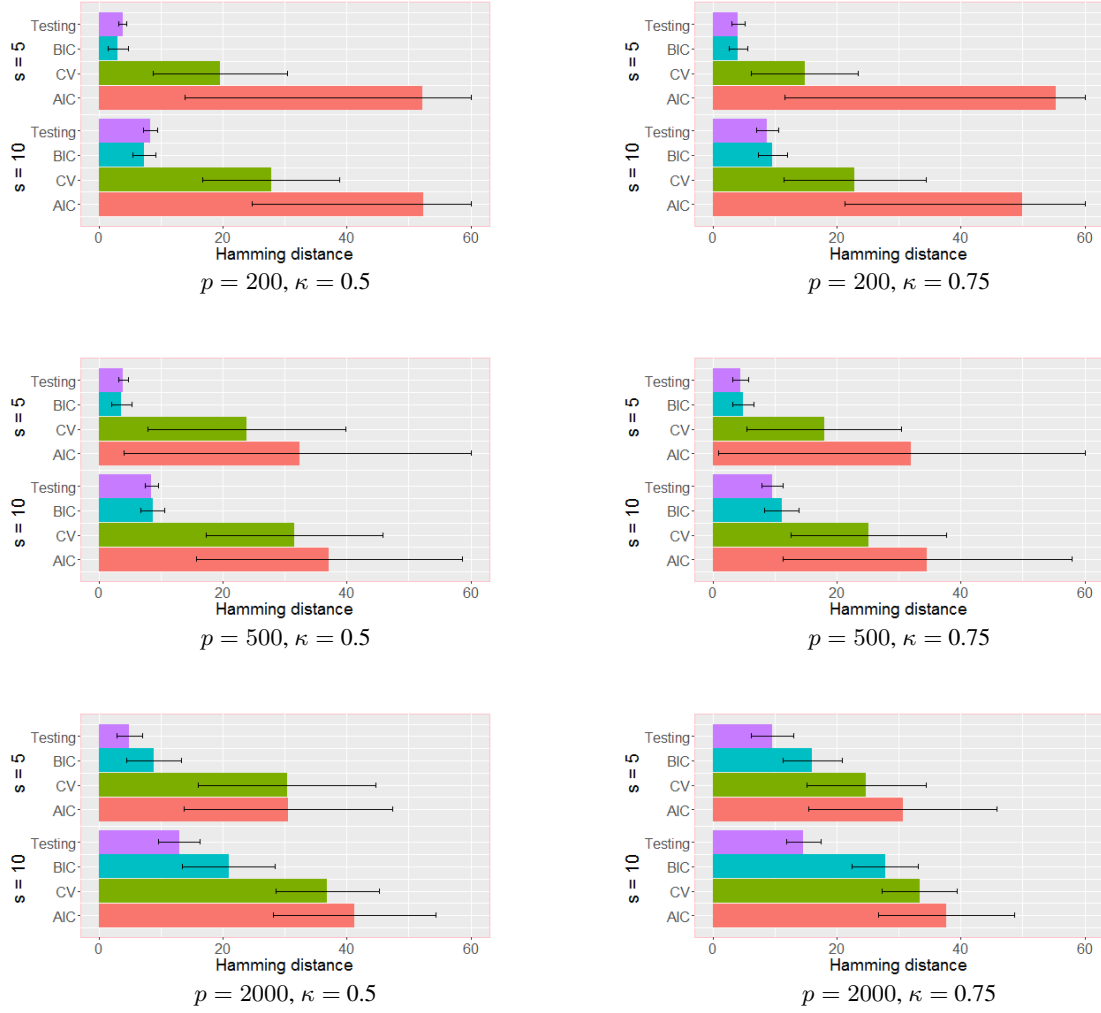


Fig. 1: Variable selection errors of ℓ_1 -regularized logistic regression with four different calibration schemes for the tuning parameter. The 12 simulation settings differ in the number of variables p , correlation κ , and sparsity s .




Fig. 2: Run times (in seconds) of ℓ_1 -regularized logistic regression with four different calibration schemes for the tuning parameter. Depicted are the results for $p = 2000$ and $\kappa \in \{0.5, 0.75\}$.

4. REAL DATA APPLICATIONS

In this section, we apply the proposed scheme to real data. We consider three data sets:

- a) Gene expression data from a leukemia microarray study (Golub et al., 1999). The data comprises $n = 72$ patients; 25 patients with acute myeloid leukemia and 47 patients with acute lymphoblastic leukemia. The predictors are the expression levels of $p = 7129$ genes. The data is summarized in the R package `golubEsets`. The goal is to select the genes whose expression levels discriminate between the two types of leukemia.
- b) The above data with the additional preprocessing and filtering described in (Dudoit et al., 2002). This reduces the number of genes to $p = 3571$. The data is summarized in the R package `cancerclass`.
- c) Proteomics data from a melanoma study (Mian et al., 2005). The data comprises $n = 205$ patients; 101 patients with stage I (moderately severe) melanoma and 104 patients with stage IV (very severe) melanoma. The raw data contains the intensities of 18'856 mass-charge (m/z) values measured in the patients' serum samples. We apply the preprocessing described in (Vasiliu et al., 2014), which results in $p = 500$ m/z values, and we subsequently normalize the data. The goal is to select the m/z values whose intensities discriminate between the two melanoma stages.

The objective of our method is feature selection. However, since there are no ground truths available for the above applications, we cannot measure feature selection accuracy directly. Instead, we need to infer the method's performance from the number of selected predictors and the prediction accuracy. We generally seek methods that yield a model with a small number of predictors (easy to interpret) and small prediction errors (good fit of the data). Moreover, an increase in prediction accuracy through refitting indicates well-estimated supports, while a deterioration in prediction accuracy through refitting indicates false negatives or false positives. We thus report the model sizes and the prediction errors of Leave-One-Out Cross-Validation without (LOOCV) and with refitting (LOOCV-refit). Typically, no method is simultaneously dominating in all measures, so that one needs to weight the two aspects according to the objective. For example, the model size is sometimes considered secondary when the goal is prediction, but it is a crucial factor for support recovery.

We apply the four different methods as described in the previous section. The results are summarized in Table 1. We observe that the methods form two clusters: On the one hand, CV and AIC provide the most accurate predictions. On the other hand, BIC and the testing-based approach select considerably smaller models and show a larger increase in accuracy after refitting. This is expected, in view of CV and AIC being designed for prediction, and BIC and testing being designed for feature selection. The clustering suggests that the testing-method should be compared especially with BIC. In the first two data examples, the testing-based method is dominating BIC, because it provides more accurate prediction with smaller models. In the third example, BIC is more accurate in prediction, but the testing-based approach provides reasonable prediction (compare especially with CV and AIC after refitting) with only one variable.

Table 1: Means and standard deviations of the model sizes and of the misclassification rates for Leave-One-Out Cross-Validation without and with refitting.

Method	Model size	LOOCV	LOOCV-refit
a) Gene expression data with $p = 7129$ genes			
Testing	4.35 (1.36)	0.153 (0.362)	0.111 (0.316)
BIC	5.03 (2.78)	0.181 (0.387)	0.111 (0.316)
CV	25.75 (3.25)	0.056 (0.231)	0.042 (0.201)
AIC	20.36 (3.07)	0.069 (0.256)	0.069 (0.256)
b) Gene expression data with $p = 3571$ genes			
Testing	4.42 (1.39)	0.167 (0.375)	0.125 (0.333)
BIC	4.99 (2.73)	0.194 (0.399)	0.139 (0.348)
CV	25.28 (2.91)	0.056 (0.231)	0.069 (0.256)
AIC	20.17 (3.41)	0.083 (0.278)	0.056 (0.231)
c) Proteomics data with $p = 500$ m/z values			
Testing	1.00 (0.00)	0.205 (0.405)	0.205 (0.405)
BIC	13.84 (1.26)	0.117 (0.322)	0.117 (0.322)
CV	23.52 (3.22)	0.117 (0.322)	0.195 (0.397)
AIC	26.86 (5.38)	0.122 (0.328)	0.185 (0.390)

5. DISCUSSION

We have introduced a scheme for the calibration of ℓ_1 -penalized likelihood for feature selection in logistic regression. A distinctive feature of the approach are its theoretical guarantees. Indeed, the new method satisfies optimal finite sample bounds, while for existing methods, the available theory is limited to asymptotic results - or there is no theory at all. Given that in applications, sample sizes are always finite, only finite sample theory can provide concrete guidance for practitioners. In addition to the theory, the scheme is easy to implement, computationally efficient, and competitive in simulations and real data applications. Our contribution thus shows that the testing ideas introduced by Chichignoud et al. (2014) are applicable well beyond linear regression.

A topic for further research are the design assumptions. The focus of this paper is feature selection, where strict assumptions on the correlations in X cannot be avoided. However, it would be interesting to extend our approach to tasks that are less sensitive to correlations, such as ℓ_2 -estimation and prediction. Also, the choice $C = 1.5$ is supported both by theory and simulations, but it can slightly disagree with the lower bound $C \geq 1.5a/c_{\min}$ if there are correlations. A further study of this lower bound might be relevant especially for tasks that allow for stronger correlations.

ACKNOWLEDGEMENT

Wei Li acknowledges support from the China Scholarship Council. We thank Haiyun Jin for his valuable contributions to the implementation of our method, and we thank Andrew Zhou for the inspiring discussions and the insightful remarks.

REFERENCES

- BÜHLMANN, P. & VAN DE GEER, S. (2011). *Statistics for High-dimensional Data: Methods, Theory and Applications*. Springer.
- BUNEA, F. (2008). Honest variable selection in linear and logistic regression models via ℓ_1 and $\ell_1 + \ell_2$ penalization. *Electron. J. Stat.* **2**, 1153–1194.
- CHICHIGNOUD, M., LEDERER, J. & WAINWRIGHT, M. (2014). Tuning lasso for sup-norm optimality. *arXiv:1410.0247*.
- DALALYAN, A., HEBIRI, M. & LEDERER, J. (2016). On the Prediction Performance of the Lasso. *To appear in Bernoulli*.
- DUDOIT, S., FRIDLYAND, J. & SPEED, T. (2002). Comparison of discrimination methods for the classification of tumors using gene expression data. *J. Am. Statist. Assoc.* **97**, 77–87.
- FAN, Y. & TANG, C. (2013). Tuning parameter selection in high dimensional penalized likelihood. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **75**, 531–552.
- FRIEDMAN, J., HASTIE, T., SIMON, N. & TIBSHIRANI, R. (2016). glmnet: Lasso and elastic-net regularized generalized linear models. *R package version 2.0-5*.
- GOLUB, T., SLONIM, D., TAMAYO, P., HUARD, C., GAASENBEEK, M., MESIROV, J., COLLIER, H., LOH, M., DOWNING, J., CALIGIURI, M., BLOOMFIELD, D. & LANDER, E. (1999). Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring. *Science* **286**, 531–537.
- HASTIE, T., TIBSHIRANI, R. & WAINWRIGHT, M. (2015). *Statistical Learning With Sparsity: The Lasso and Generalizations*. CRC Press.
- LOUNICI, K. (2008). Sup-norm convergence rate and sign concentration property of Lasso and Dantzig estimators. *Electron. J. Stat.* **2**, 90–102.
- MIAN, S., UGUREL, S., PARKINSON, E., SCHLENZKA, I., DRYDEN, I., LANCASHIRE, L., BALL, G., CREASER, C., REES, R. & SCHADENDORF, D. (2005). Serum proteomic fingerprinting discriminates between clinical stages and predicts disease progression in melanoma patients. *J. Clin. Oncol.* **23**, 5088–5093.
- R CORE TEAM (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>.
- RAVIKUMAR, P., WAINWRIGHT, M. & LAFFERTY, J. (2010). High-dimensional Ising model selection using ℓ_1 -regularized logistic regression. *Ann. Statist.* **38**, 1287–1319.
- RYALI, S., SUPEKAR, K., ABRAMS, D. & MENON, V. (2010). Sparse logistic regression for whole-brain classification of fmri data. *Neuroimage* **51**, 752–764.
- SABOURIN, J., VALDAR, W. & NOBEL, A. (2015). A permutation approach for selecting the penalty parameter in penalized model selection. *Biometrics* **71**, 1185–1194.
- SHAO, J. (1993). Linear model selection by cross-validation. *J. Am. Statist. Assoc.* **88**, 486–494.
- VASILIU, D., DEY, T. & DRYDEN, I. (2014). Penalized Euclidean Distance Regression. *arXiv:1405.4578*.
- WU, T., CHEN, Y., HASTIE, T., SOBEL, E. & LANGE, K. (2009). Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics* **25**, 714–721.
- ZHANG, Y., LI, R. & TSAI, C.-L. (2010). Regularization parameter selections via generalized information criterion. *J. Amer. Statist. Assoc.* **105**, 312–323.
- ZHAO, P. & YU, B. (2006). On model selection consistency of lasso. *J. Mach. Learn. Res.* **7**, 2541–2563.

APPENDIX

In this appendix, we provide proofs for our theoretical claims, and we present additional simulation results.

Proofs for the Theoretical Claims

We provide here proofs for Theorems 1 and 2. Throughout this section, we write $\rho(u, v) = \exp(u^\top v) / (1 + \exp(u^\top v))$ for vectors u, v of the same length. For ease of notation, we will suppress the subscript λ at most instances.

The key quantities in the proofs are two vectors $\hat{\alpha} = (\hat{\alpha}_S^\top, \hat{\alpha}_{S^c}^\top)^\top \in \mathbb{R}^p$ and $\hat{\nu} = (\hat{\nu}_S^\top, \hat{\nu}_{S^c}^\top)^\top \in \mathbb{R}^p$ constructed as follows:

1. define the primal subvector $\hat{\alpha}_S \in \mathbb{R}^s$ such that

$$\hat{\alpha}_S \in \operatorname{argmin}_{\theta \in \mathbb{R}^s} \left\{ \sum_{i=1}^n (\log(1 + \exp(x_{i,S}^\top \theta)) - y_i x_{i,S}^\top \theta) / n + \lambda \|\theta\|_1 \right\};$$

2. set $\hat{\alpha}_{S^c} = 0 \in \mathbb{R}^{p-s}$;
3. define the dual vector $\hat{\nu} \in \mathbb{R}^p$ via its elements

$$\hat{\nu}_j = \sum_{i=1}^n X_{ij} (y_i - \rho(x_i, \hat{\alpha})) / (n\lambda) \quad (j \in [p]).$$

The proofs of the theorems are based on three auxiliary lemmas. Figure 3 depicts the dependencies.




Fig. 3: Dependencies among the lemmas and theorems. For example, the arrow between Lemmas A1 and A2 indicates that Lemma A2 relies on Lemma A1.

LEMMA A1 (ℓ_2 -BOUND FOR THE PRIMAL SUBVECTOR). *If $\lambda \leq c_{\min}^2 / (10c_b s c_{\max})$ and $\|X^\top \varepsilon / n\|_\infty \leq \lambda/4$, then*

$$\|\hat{\alpha}_S - \beta_S^*\|_2 \leq \frac{5\lambda\sqrt{s}}{c_{\min}}.$$

The proof of this lemma follows along the same lines as the proof of Lemma 3 in (Ravikumar et al., 2010).

Lemma A1 implies that the primal subvector $\hat{\alpha}_S$ is ℓ_2 -consistent, which enables us to develop a Taylor series expansion of $\rho(x_{i,S}, \hat{\alpha}_S)$ at β_S^* according to

$$\rho(x_{i,S}, \hat{\alpha}_S) - \rho(x_{i,S}, \beta_S^*) = w(x_{i,S}, \beta_S^*) x_{i,S}^\top (\hat{\alpha}_S - \beta_S^*) + r_i \quad (\text{A1})$$

with remainder term

$$r_i = (\hat{\alpha}_S - \beta_S^*)^\top \int_0^1 (\nabla^2 \rho(x_{i,S}, \hat{\alpha}_S + t(\hat{\alpha}_S - \beta_S^*))) (1-t) dt (\hat{\alpha}_S - \beta_S^*).$$

The derivative reads explicitly

$$\nabla^2 \rho(x_{i,S}, \hat{\alpha}_S + t(\hat{\alpha}_S - \beta_S^*)) = \xi_i(t) x_{i,S} x_{i,S}^\top, \quad (\text{A2})$$

where $\xi_i(t) = \exp(\eta_i(t))(1 - \exp(\eta_i(t)))/(1 + \exp(\eta_i(t)))^3$ and $\eta_i(t) = x_{i,S}^\top(\hat{\alpha}_S + t(\hat{\alpha}_S - \beta_S^*))$ for $t \in [0, 1]$ and $i \in [n]$. Summarizing r_1, \dots, r_n from Equation (A1) in the vector $r = (r_1, \dots, r_n)^\top$, we can now state the following result.

LEMMA A2 (ℓ_∞ -BOUND FOR THE REMAINDER TERM). *If $\lambda \leq \gamma c_{\min}^2 / (100c_b(2 - \gamma)sc_{\max})$ and $\|X^\top \varepsilon / n\|_\infty \leq \lambda/4$, then*

$$\|X^\top r / n\|_\infty \leq \frac{\lambda\gamma}{4(2 - \gamma)}.$$

Proof. The proof follows readily from Lemma A1. To see this, note that because $|X_{ij}| \leq c_b$ for all $i \in [n]$ and $j \in [p]$, it holds that

$$|r^\top x^j / n| = \left| \sum_{i=1}^n X_{ij} r_i / n \right| \leq \sum_{i=1}^n |X_{ij}| |r_i| / n \leq \sum_{i=1}^n c_b |r_i| / n$$

for all $j \in [p]$. By the closed form of r_i in Equation (A2), it holds that $|\xi_i(t)| \leq 1$ for all $t \in [0, 1]$, and since $\hat{\alpha}_{S^c} = \beta_{S^c}^* = 0$, we then get

$$\begin{aligned} |r^\top x^j / n| &\leq c_b(\hat{\alpha}_S - \beta_S^*)^\top \left(\sum_{i=1}^n x_{i,S} x_{i,S}^\top / n \right) (\hat{\alpha}_S - \beta_S^*) \\ &= c_b(\hat{\alpha}_S - \beta_S^*)^\top (X_S^\top X_S / n) (\hat{\alpha}_S - \beta_S^*) \\ &= c_b(\hat{\alpha} - \beta^*)^\top (X^\top X / n) (\hat{\alpha} - \beta^*) \\ &\leq c_b c_{\max} \|\hat{\alpha}_S - \beta_S^*\|_2^2. \end{aligned} \quad (\text{A3})$$

Moreover, because $\lambda \leq \gamma c_{\min}^2 / (100c_b(2 - \gamma)sc_{\max}) \leq c_{\min}^2 / (10c_b sc_{\max})$ and $\|X^\top \varepsilon / n\|_\infty \leq \lambda/4$, the assumptions of Lemma A1 are satisfied. Combining this lemma with Equation (A3) yields

$$|r^\top x^j / n| \leq \frac{25c_b \lambda^2 sc_{\max}}{c_{\min}^2} \leq \frac{\lambda\gamma}{4(2 - \gamma)}$$

for each $j \in [p]$. Thus, $\|X^\top r / n\|_\infty \leq \lambda\gamma / (4(2 - \gamma))$ as desired. \square

LEMMA A3 (PRIMAL DUAL WITNESS CONSTRUCTION). *The pair $(\hat{\alpha}, \hat{\nu})$ defined above satisfies the following three properties:*

- (i) *It holds that $\hat{\nu}_S \in \partial \|\hat{\alpha}_S\|_1$;*
- (ii) *If $\|\hat{\nu}_{S^c}\|_\infty < 1$, then any solution $\hat{\beta}$ to the problem (2) satisfies $\text{supp}(\hat{\beta}) \subset S$;*
- (iii) *Under Assumption 1 and $\|\hat{\nu}_{S^c}\|_\infty < 1$, the solution $\hat{\beta}$ is unique, and $\hat{\beta} = \hat{\alpha} = (\hat{\alpha}_S^\top, \hat{\alpha}_{S^c}^\top)^\top \in \mathbb{R}^p$.*

Proof. We conduct the proof in three steps in correspondence with the three claims.

Step 1: We show that if $\|\hat{\nu}_{S^c}\|_\infty \leq 1$, the pair $(\hat{\alpha}, \hat{\nu})$ satisfies the KKT conditions, that is, $\hat{\nu} \in \partial \|\hat{\alpha}\|_1$ and

$$-\sum_{i=1}^n X_{ij} (y_i - \rho(x_i, \hat{\alpha})) / n + \lambda \hat{\nu}_j = 0$$

for $j \in [p]$. By 1. in the construction at the beginning of this section, there is a $\kappa \in \partial \|\hat{\alpha}_S\|_1 \subset \mathbb{R}^s$ such that

$$-\sum_{i=1}^n X_{ij} (y_i - \rho(x_{i,S}, \hat{\alpha}_S)) / n + \lambda \kappa_j = 0$$

for $j \in S$. Hence, with $\hat{\alpha}_{S^c} = 0$ in 2. and the definition of $\hat{\nu}_j$ in 3.,

$$\hat{\nu}_j = \sum_{i=1}^n X_{ij} (y_i - \rho(x_{i,S}, \hat{\alpha}_S)) / (n\lambda) = \kappa_j$$

for $j \in S$, that is, $\hat{\nu}_S \in \partial \|\hat{\alpha}_S\|_1$ as desired.

Step 2: We now show that $\text{supp}(\hat{\beta}) \subset S$ for all $\hat{\beta} \in \mathbb{R}^p$ that satisfy

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \{L(\beta) + \lambda \|\beta\|_1\}.$$

In view of the condition $\|\hat{\nu}_{S^c}\|_\infty < 1$ and of Step 1, the pair $(\hat{\alpha}, \hat{\nu})$ satisfies the KKT conditions for the above problem and thus, $\hat{\alpha}$ is a minimizer of the objective function. Consequently,

$$L(\hat{\beta}) + \lambda \|\hat{\beta}\|_1 = L(\hat{\alpha}) + \lambda \|\hat{\alpha}\|_1.$$

Since, $\hat{\nu} \in \partial \|\hat{\alpha}\|_1$ by Step 1, it holds that $\|\hat{\alpha}\|_1 = \langle \hat{\nu}, \hat{\alpha} \rangle$. Plugging this into the previous display yields

$$L(\hat{\beta}) + \lambda \|\hat{\beta}\|_1 = L(\hat{\alpha}) + \lambda \langle \hat{\nu}, \hat{\alpha} \rangle.$$

We can now subtract $\lambda \langle \hat{\nu}, \hat{\beta} \rangle$ on both sides to obtain

$$L(\hat{\beta}) + \lambda \|\hat{\beta}\|_1 - \lambda \langle \hat{\nu}, \hat{\beta} \rangle = L(\hat{\alpha}) + \lambda \langle \hat{\nu}, \hat{\alpha} - \hat{\beta} \rangle.$$

By 2. and 3. in the above construction, it holds that $\lambda \hat{\nu} = -L'(\hat{\alpha})$, where $L'(\cdot)$ denotes the derivative of $L(\cdot)$. Thus, we can further deduce

$$\lambda \|\hat{\beta}\|_1 - \lambda \langle \hat{\nu}, \hat{\beta} \rangle = L(\hat{\alpha}) - \langle L'(\hat{\alpha}), \hat{\alpha} - \hat{\beta} \rangle - L(\hat{\beta}).$$

Because the Hessian of $L(\beta)$ is a non-negative matrix, $L(\cdot)$ is a convex function. It holds that

$$L(\hat{\beta}) \geq L(\hat{\alpha}) + \langle L'(\hat{\alpha}), \hat{\beta} - \hat{\alpha} \rangle.$$

Combining the two displays yields

$$\lambda \|\hat{\beta}\|_1 - \lambda \langle \hat{\nu}, \hat{\beta} \rangle \leq 0,$$

and dividing by the tuning parameter yields further

$$\|\hat{\beta}\|_1 \leq \langle \hat{\nu}, \hat{\beta} \rangle.$$

However, by Hölder's inequality and $\|\hat{\nu}\|_\infty \leq 1$, it holds that

$$\|\hat{\beta}\|_1 \geq \langle \hat{\nu}, \hat{\beta} \rangle.$$

Consequently,

$$\|\hat{\beta}\|_1 = \langle \hat{\nu}, \hat{\beta} \rangle.$$

In view of the condition $\|\hat{\nu}_{S^c}\|_\infty < 1$, this can only be true if $\hat{\beta}_j = 0$ for all $j \in S^c$. This completes the proof of Step 2.

Step 3: We now show that $\hat{\beta} = \hat{\alpha}$. From Step 2, we deduce that $\hat{\beta} = (\hat{\beta}_S^\top, 0)^\top$ with

$$\hat{\beta}_S \in \arg \min_{\theta \in \mathbb{R}^s} \left\{ \sum_{i=1}^n (\log(1 + \exp(x_{i,S}^\top \theta)) - y_i x_{i,S}^\top \theta) / n + \lambda \|\theta\|_1 \right\}.$$

Moreover, since the minimal eigenvalue of $X_S^\top W X_S/n$ is larger than zero by Assumption 1, this problem has a unique solution. Combining this with 1. in the construction at the beginning of this section yields $\hat{\beta}_S = \hat{\alpha}_S$, that is, $\hat{\beta} = \hat{\alpha}$. \square

Proof of Theorem 1. We conduct the proof in two steps. The first step is to show that $\text{supp}(\hat{\beta}) \subset S$ and that $\hat{\beta}$ is the unique solution of the problem (2). The second step is to show the ℓ_∞ -bound and the result on support recovery.

Step 1: We first show that $\text{supp}(\hat{\beta}) \subset S$ and that $\hat{\beta}$ is the unique solution. This result holds true if the primal-dual pair $(\hat{\alpha}, \hat{\nu}) \in \mathbb{R}^p \times \mathbb{R}^p$ constructed as in Lemma A3 satisfies $\|\hat{\nu}_{S^c}\|_\infty < 1$. To show the latter inequality, we use the definition of ε and Equation (A1) to rewrite 3. in the construction above as

$$\sum_{i=1}^n X_{ij} w(x_{i,S}, \beta_S^*) x_{i,S}^\top (\hat{\alpha}_S - \beta_S^*)/n - \sum_{i=1}^n X_{ij} (\varepsilon_i - r_i)/n + \lambda \hat{\nu}_j = 0 \quad (j \in [p]).$$

Because $w(x_{i,S}, \beta_S^*) = w(x_i, \beta^*)$ for each $i \in [n]$, we can put the above display in the matrix form

$$(X^\top W X/n) \begin{pmatrix} \beta_S^* - \hat{\alpha}_S \\ 0 \end{pmatrix} + X^\top (\varepsilon - r)/n - \lambda \begin{pmatrix} \hat{\nu}_S \\ \hat{\nu}_{S^c} \end{pmatrix} = 0,$$

and then in the block matrix form

$$n^{-1} \begin{pmatrix} X_S^\top W X_S & X_S^\top W X_{S^c} \\ X_{S^c}^\top W X_S & X_{S^c}^\top W X_{S^c} \end{pmatrix} \begin{pmatrix} \beta_S^* - \hat{\alpha}_S \\ 0 \end{pmatrix} + n^{-1} \begin{pmatrix} X_S^\top (\varepsilon - r) \\ X_{S^c}^\top (\varepsilon - r) \end{pmatrix} - \lambda \begin{pmatrix} \hat{\nu}_S \\ \hat{\nu}_{S^c} \end{pmatrix} = 0.$$

We now solve this equation for $\lambda \hat{\nu}_{S^c}$ and find

$$\lambda \hat{\nu}_{S^c} = X_{S^c}^\top W X_S (\beta_S^* - \hat{\alpha}_S)/n + X_{S^c}^\top (\varepsilon - r)/n.$$

Since the matrix $X_S^\top W X_S$ is invertible by Assumption 2, we can solve the block matrix equation also for $(\beta_S^* - \hat{\alpha}_S)/n$ and find

$$(\beta_S^* - \hat{\alpha}_S)/n = -(X_S^\top W X_S)^{-1} X_S^\top (\varepsilon - r)/n + \lambda (X_S^\top W X_S)^{-1} \hat{\nu}_S. \quad (\text{A4})$$

Combining the two displays yields

$$\lambda \hat{\nu}_{S^c} = -X_{S^c}^\top W X_S (X_S^\top W X_S)^{-1} X_S^\top (\varepsilon - r)/n + X_{S^c}^\top (\varepsilon - r)/n + \lambda X_{S^c}^\top W X_S (X_S^\top W X_S)^{-1} \hat{\nu}_S. \quad (\text{A5})$$

Taking ℓ_∞ -norms on both sides of Equation (A5) and using the triangle inequality, we find

$$\begin{aligned} \|\lambda \hat{\nu}_{S^c}\|_\infty &\leq \|X_{S^c}^\top W X_S (X_S^\top W X_S)^{-1} X_S^\top (\varepsilon - r)/n\|_\infty + \|X_{S^c}^\top (\varepsilon - r)/n\|_\infty \\ &\quad + \lambda \|X_{S^c}^\top W X_S (X_S^\top W X_S)^{-1} \hat{\nu}_S\|_\infty. \end{aligned}$$

Invoking properties of the induced matrix norms and the ℓ_∞ -norm and the condition $\|\hat{\nu}_S\|_\infty \leq 1$ deduced in Lemma A3, and rearranging the terms then provide us with

$$\|\lambda \hat{\nu}_{S^c}\|_\infty \leq \|X_{S^c}^\top W X_S (X_S^\top W X_S)^{-1}\|_\infty (\|X^\top (\varepsilon - r)/n\|_\infty + \lambda) + \|X^\top (\varepsilon - r)/n\|_\infty.$$

Next, we divide by λ on both sides, apply Assumption 2, use the triangle inequality, and rearrange the terms again to find

$$\|\hat{\nu}_{S^c}\|_\infty \leq (1 - \gamma) + \frac{2 - \gamma}{\lambda} (\|X^\top \varepsilon/n\|_\infty + \|X^\top r/n\|_\infty).$$

By the definition of \mathcal{T}_λ , it holds that $\|X^\top \varepsilon/n\|_\infty \leq \lambda \gamma / (4(2 - \gamma))$, which is equivalent to

$$\frac{2 - \gamma}{\lambda} \|X^\top \varepsilon/n\|_\infty \leq \frac{\gamma}{4}.$$

Since $\gamma \in (0, 1]$, the condition $\|X^\top \varepsilon/n\|_\infty \leq \gamma/4$ in Lemma A2 is satisfied on the event \mathcal{T}_λ . Combining this with the assumption $\lambda \leq \gamma c_{\min}^2 / (100 c_b (2 - \gamma) s c_{\max})$ implies that $\|X^\top r/n\|_\infty \leq \lambda \gamma / (4(2 - \gamma))$,

see Lemma A2. Thus,

$$\|\hat{\nu}_{S^c}\|_\infty \leq (1 - \gamma) + \frac{\gamma}{4} + \frac{2 - \gamma}{\lambda} \|X^\top r/n\|_\infty \leq (1 - \gamma) + \frac{\gamma}{4} + \frac{\gamma}{4} < 1.$$

We finally invoke Lemma A3 to conclude that $\text{supp}(\hat{\beta}) \subset S$ and that $\hat{\beta} = \hat{\alpha} = (\hat{\alpha}_S^\top, 0^\top)^\top \in \mathbb{R}^p$ is the unique solution of the problem (2), as desired.

Step 2: To show the ℓ_∞ -bound, we use Equation (A4) and $\hat{\beta}_S = \hat{\alpha}_S$ from Step 1 and find

$$\begin{aligned} \beta_S^* - \hat{\beta}_S &= -(X_S^\top W X_S)^{-1} X_S^\top (\varepsilon - r) + \lambda n (X_S^\top W X_S)^{-1} \hat{\nu}_S \\ &= -(X_S^\top W X_S/n)^{-1} (X_S^\top (\varepsilon - r)/n) + \lambda (X_S^\top W X_S/n)^{-1} \hat{\nu}_S. \end{aligned}$$

We then find similarly as before

$$\|\beta_S^* - \hat{\beta}_S\|_\infty \leq \|(X_S^\top W X_S/n)^{-1}\|_\infty (\lambda + \|X^\top (\varepsilon - r)/n\|_\infty).$$

By the definition of a , we have

$$\|(X_S^\top W X_S/n)^{-1}\|_\infty \leq a \|(X_S^\top W X_S/n)^{-1}\|_2 \leq a/c_{\min}.$$

Combining this with the bounds on $\|X^\top \varepsilon/n\|_\infty$ and $\|X^\top r/n\|_\infty$ deduced in Step 1 yields

$$\|\beta_S^* - \hat{\beta}_S\|_\infty \leq \frac{a}{c_{\min}} \left(\lambda + \frac{\lambda\gamma}{4(2-\gamma)} + \frac{\lambda\gamma}{4(2-\gamma)} \right) \leq 1.5a\lambda/c_{\min}.$$

Since $\text{supp}(\hat{\beta}) \subset S$ by Step 1, the above display implies that

$$\|\beta^* - \hat{\beta}\|_\infty \leq 1.5a\lambda/c_{\min}.$$

Consequently, $\text{supp}(\hat{\beta}) = S$ as long as $\min_{j \in S} |\beta_j^*| > 1.5a\lambda/c_{\min}$. This concludes the proof. \square

Proof of Theorem 2. The proof is conducted in three steps. The first step is to show the bound on $\hat{\lambda}$, the second step is to show the bound on the sup-norm error, and the last step is to show that $\hat{S} \supset S$. To begin with, we define the event

$$\mathcal{T}_\delta^* = \left\{ \frac{4(2-\gamma)}{n\gamma} \|X^\top \varepsilon\|_\infty \leq \lambda_\delta^* \right\}.$$

By our definition of the oracle tuning parameter in (4), we have that $\text{pr}(\mathcal{T}_\delta^*) \geq 1 - \delta$. Thus, it suffices to show that the results hold conditioned on the event \mathcal{T}_δ^* .

Step 1: To show that $\hat{\lambda} \leq \lambda_\delta^*$, we proceed by proof by contradiction. If $\hat{\lambda} > \lambda_\delta^*$, then the definition of our testing-based calibration implies that there must exist two tuning parameters $\lambda', \lambda'' \geq \lambda_\delta^*$ such that

$$\|\hat{\beta}_{\lambda'} - \hat{\beta}_{\lambda''}\|_\infty > C(\lambda' + \lambda''). \quad (\text{A6})$$

However, because both $\mathcal{T}_{\lambda'}$ and $\mathcal{T}_{\lambda''}$ include \mathcal{T}_δ^* , and because $C \geq 1.5a/c_{\min}$, Theorem 1 implies that $\|\hat{\beta}_{\lambda'} - \beta^*\|_\infty \leq C\lambda'$ and $\|\hat{\beta}_{\lambda''} - \beta^*\|_\infty \leq C\lambda''$. By applying the triangle inequality, we have

$$\|\hat{\beta}_{\lambda'} - \hat{\beta}_{\lambda''}\|_\infty \leq \|\hat{\beta}_{\lambda'} - \beta^*\|_\infty + \|\hat{\beta}_{\lambda''} - \beta^*\|_\infty \leq C(\lambda' + \lambda'').$$

This upper bound contradicts our earlier conclusion (A6) and, therefore, yields the desired bound on the tuning parameter.

Step 2: On the event \mathcal{T}_δ^* , we have $\hat{\lambda} \leq \lambda_\delta^*$, and so the testing-based method implies that

$$\|\hat{\beta}_{\hat{\lambda}} - \hat{\beta}_{\lambda_\delta^*}\|_\infty \leq C(\hat{\lambda} + \lambda_\delta^*) \leq 2C\lambda_\delta^*.$$

By applying the triangle inequality, we find that

$$\|\hat{\beta}_{\hat{\lambda}} - \beta^*\|_{\infty} \leq \|\hat{\beta}_{\hat{\lambda}} - \hat{\beta}_{\lambda_{\delta}^*}\|_{\infty} + \|\hat{\beta}_{\lambda_{\delta}^*} - \beta^*\|_{\infty} \leq 2C\lambda_{\delta}^* + \|\hat{\beta}_{\lambda_{\delta}^*} - \beta^*\|_{\infty}.$$

Theorem 1 implies that $\|\hat{\beta}_{\lambda_{\delta}^*} - \beta^*\|_{\infty} \leq 1.5a\lambda_{\delta}^*/c_{\min} \leq C\lambda_{\delta}^*$, and combining the pieces yields the desired sup-norm bound.

Step 3: Let us finally show that $\hat{S} \supset S$. Suppose $j \in S$, then by the bound on the sup-norm error that we deduce in Step 2, we have

$$|(\hat{\beta}_{\hat{\lambda}})_j| \geq |\beta_j^*| - 3C\lambda_{\delta}^*.$$

In view of the condition $\min_{j \in S} |\beta_j^*| > 6C\lambda_{\delta}^*$ and the definition $\hat{S} = \{j \in [p] : |(\hat{\beta}_{\hat{\lambda}})_j| \geq 3C\hat{\lambda}\}$, we conclude that $j \in \hat{S}$, that is, $\hat{S} \supset S$. This completes the proof. \square

Additional Simulations

We finally present additional simulation results. The settings are described in Section 3. The simulation parameters here are, according to our experience, particularly favorable for BIC. However, we observe that the testing-based scheme still rivals BIC (and outmatches all other methods) in feature selection and is faster in computations. This corroborates that beyond its main advantage, the theoretical guarantees, the proposed scheme is also a competitor in practice.




Fig. 4: Variable selection errors of ℓ_1 -regularized logistic regression with four different calibration schemes for the tuning parameter. The 12 simulation settings differ in the number of variables p , correlation κ , and sparsity s .




Fig. 5: Run times (in seconds) of ℓ_1 -regularized logistic regression with four different calibration schemes for the tuning parameter. Depicted are the results for $p = 2000$ and $\kappa \in \{0, 0.25\}$.